arXiv:1403.0783v1 [cs.DB] 4 Mar 2014

Uncertainty in Crowd Data Sourcing under Structural Constraints

Antoine Amarilli[1], Yael Amsterdamer[2], and Tova Milo[2]

[1] Institut Mines–Télécom; Télécom ParisTech; CNRS LTCI, Paris, France
[2] Tel Aviv University, Tel Aviv, Israel

Abstract. Applications extracting data from crowdsourcing platforms must deal with the *uncertainty* of crowd answers in two different ways: first, by deriving estimates of the correct value from the answers; second, by choosing crowd questions whose answers are expected to minimize this uncertainty relative to the overall data collection goal. Such problems are already challenging when we assume that questions are unrelated and answers are independent, but they are even more complicated when we assume that the unknown values follow hard *structural constraints* (such as monotonicity).

In this vision paper, we examine how to formally address this issue with an approach inspired by [2]. We describe a generalized setting where we model constraints as linear inequalities, and use them to guide the choice of crowd questions and the processing of answers. We present the main challenges arising in this setting, and propose directions to solve them.

1 Introduction

Crowd data sourcing leverages human knowledge to obtain information which does not exist in conventional databases. This may be done by posing targeted questions to crowd users, through conventional crowdsourcing platforms such as Amazon Mechanical Turk [6]. Contrary to many works that use the crowd as a *means* to perform different tasks, here the crowd serves as a *source* of information.

Many challenges arise when using the crowd as a data source. First, human answers have a high latency and are usually provided against some (monetary) compensation, so we must minimize the number of posed questions. Second, answers collected from the crowd may be erroneous and noisy, so we must control and improve answer quality, e.g., pose the same question to multiple workers.

A vast body of research has tackled these issues for various data procurement tasks (e.g., [1,2,7,8,11,12]). For example, [7] studied the number of answers that must be obtained to reach sufficient confidence in the final answer of a given Boolean question, and mentions the problem of deciding, when there are several questions to answer, which is the *next best question* to ask the crowd. In different situations [2,12], this selection of questions is performed by comparing the expected contribution of the answers to some data acquisition goal. However, in such situations, the answers to the various questions are independent, so that we can choose the next best question by looking at each question in isolation.

In this paper, we study the problem of collecting numerical values from the crowd under *hard a-priori constraints* on the final answers, caused by inherent data dependencies.

For instance, suppose that we have devised a lossy compression algorithm for e.g. music files, and that we wish to estimate the average quality rating of different compression ratios in a user population. We can ask a few random crowd workers to evaluate the quality $q_1, \ldots, q_n$ for each of n compression ratios of *increasing* lossiness. The quality ratings of any given person are not independent: we can assume that every person will consider q_1 (the quality of the least lossy compression) to be at least as high as q_2, and so on. Consequently, the average q_1 in the entire population is higher than q_2, and so on. However, the quality q_1 for *some* people might be lower than q_2 for others; hence, by asking random workers we may obtain an estimation of the quality ratings that is not perfectly monotone. This use case will be our running example throughout the paper.

As another example, consider the estimation of the price that people are willing to pay for varying combined deals. In fields such as auction study in game theory [9], it is customary to assume that the price function for each user is monotone, i.e., adding products cannot decrease the deal price. For instance, we know that in the entire population, the average value of a flight and hotel cannot be lower than that of the flight alone. But again, if we sample different users for each deal, we may obtain a non-monotone estimation for the average price.

A similar problem occurs in [2], where the crowd is used to estimate the frequency of patterns in user habits. While these frequencies are dependent for patterns with overlapping activities (e.g., if someone never swims, they also never swim and dive), such dependencies are not accounted for in [2]. In general, existing work on crowd data sourcing has mostly ignored the problem of uncertainty when dealing with dependent questions [1,8]. There are works that deal in a non-trivial way with the interaction between uncertainty and dependency [4,10], but they assume that the individual outcomes observed are Boolean and not numeric like in the present paper.

We consider here two important problems that arise in the context of dependent crowd questions. First, can we *improve the variable estimation* by taking dependencies into account? For instance, in our running example, if we estimate that the quality q_1 is lower than q_2 (which contradicts our monotonicity assumption), we may attempt to correct our estimation by increasing q_1 and/or decreasing q_2; or, to begin with, we can only consider estimations that comply with our monotonicity requirement. What is the right way to enforce this monotonicity, and how does it increase the quality of our estimation?

Second, we use the dependencies to *reduce the number of questions* posed to the crowd. For instance, if we estimate that the average quality rating q_1 and q_4 of the compression ratios 1 and 4 are both 6 out of 10, we do not need to ask people about q_2 and q_3. Or, as another example, if we wish to find the lossiest compression with rating at least 6 out of 10, and we estimate that q_5 and q_{10} are 8 and 3 respectively, we can interpolate $q_6, \ldots, q_9$ (using monotonicity) and estimate that the rating with value closest to 6 is most likely q_7.

Paper structure. We first give a formal definition of the considered problem in Section 2. We next present in Section 3 a general scheme to solve the problem in the absence of dependencies, and turn in Section 4 to how dependencies should be handled. For numerous dependent variables, interpolating samples is crucial: we discuss this in Section 5. Last, we conclude in Section 6.

2 Problem statement

We wish to learn n numerical values $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)$ from the crowd. We model the *distribution of crowd answers* to questions about these values using n *random variables* $X_1, \ldots, X_n$. We assume that the *mean* of X_i is μ_i for every i.[3]

We further assume that $\boldsymbol{\mu}$ satisfies a certain known set of *linear inequalities*, represented as a matrix E of reals such that $E \cdot \boldsymbol{\mu} \leq (0)$, where (0) is the zero vector and $\cdot$ denotes the product of matrix E and vector $\boldsymbol{\mu}$. We assume that the inequalities E are feasible, namely, that there is some vector $\mathbf{e}$ satisfying E.

Example 1. In our running example, the random variables $Q_1, \ldots, Q_n$, with unknown means $q_1, \ldots, q_n$, denote the ratings obtained for the compression ratios. The inequalities represent a decreasing order: $q_2 - q_1 \leq 0$, $q_3 - q_2 \leq 0$, etc.

We consider a known *loss function* $L_{\boldsymbol{\mu}}$ which associates to a prediction $\mathbf{v}$ for the unknown values $\boldsymbol{\mu}$ some nonnegative value $L_{\boldsymbol{\mu}}(\mathbf{v})$. We assume that $L_{\boldsymbol{\mu}}$ can be written as the sum of nonnegative functions $L^i_{\mu_i}$, that is, the error function for all values is the sum of the errors of individual values. We require that for all i, $L^i_{\mu_i}(\mu_i) = 0$ and $L^i_{\mu_i}(x) \leq L^i_{\mu_i}(y)$ for all $\mu_i \leq x \leq y$ and $y \leq x \leq \mu_i$. (In other words, the loss is 0 for the correct value, and increases with the absolute error.)

Example 2. The loss function depends on the target application. For compression ratios, if our task is to find which is the lossiest compression with rating at least 6, a reasonable loss function for all variables is the *threshold* loss $L_{\mu,\tau}$ with $\tau = 6$. The value $L_{\mu,\tau}(x)$ is defined to be 1 if the x and μ are miscategorized with respect to threshold τ (formally, $\mu < \tau < x$ or $x < \tau < \mu$) and 0 otherwise. The overall loss function is the sum of the $L_{q_i,\tau}$ which counts the number of ratios that are miscategorized with respect to the threshold $\tau = 6$.

For any i, we can obtain a sample of variable X_i (we say that we *sample* X_i or *draw* X_i) by asking the corresponding question to a random crowd worker; we assume that all draws are independent both between variables and between two draws of the same variable. Our goal is to choose draws carefully and, based on the obtained samples, try to provide a prediction $\mathbf{v}$ which minimizes $L_{\boldsymbol{\mu}}(\mathbf{v})$: we phrase this in a *fixed-budget* formulation, namely minimize $L_{\boldsymbol{\mu}}(\mathbf{v})$ in expectation after a fixed number of samples.

Example 3. In the running example, sampling the variable Q_i is achieved by providing a random crowd user with a sound sample compressed with ratio i and asking for a rating for this sample. The overall objective is to choose the

[3] This assumption holds when we are interested in the *average* crowd answer, e.g., the average rating for a compression quality; and in the many cases where the errors of worker answers tend to cancel out so that the average is close to the truth [2].

right ratios for which to request more ratings, in order to minimize the number of average quality ratings that are miscategorized with respect to $\tau = 6$.

We next review the problem of minimizing the loss by choosing the "right" questions. We first study an approach for a simplified setting where there are no order constraints on the estimated values, before we consider the general case.

3 Without order constraints

Let us present a general scheme inspired by [2] for the case with no order constraints, before we extend it to order constraints in the next section.

With no constraints, as the variables are independent and the loss is the sum of the individual losses of variables, our goal is to find which one of the variables is such that one more sample for it would yield the largest loss reduction. Hence, we first focus on an individual variable X_i to describe how we predict its mean value v_i from the samples S_i observed for this variable, and how we estimate the loss reduction that we may achieve by taking one more sample.

Estimating the parameter. Our approach for a variable X given a set S of samples of this variable is to fit a model for X from the family of normal distributions, as they are a simple and general way to represent real-life data. Denote by $\Theta = \mathbb{R} \times \mathbb{R}_+$ the *parameter space*, such that every $\theta \in \Theta$, with $\theta = (\mu, \sigma^2)$, represents the normal distribution $\mathcal{N}(\mu, \sigma^2)$ with mean μ and variance σ^2. Denote by $\Pr_\theta$ the probability density function of this distribution.

As the samples S of X are assumed to be independent, we can define the probability of S according to $\mathcal{N}(\theta)$ as the product of $\Pr_\theta(s_i)$ for all $s_i \in S$. The *likelihood function* $\mathcal{L}_S : \Theta \to [0, 1]$ is then simply defined as $\mathcal{L}_S(\theta) = \Pr_\theta(S)$: it describes, *as a function of* θ, the probability[4] of the sample under θ.

Our way to fit a normal distribution to the random variable X is then the standard method of choosing the *maximum likelihood estimator* (MLE):

$$\widehat{\theta} = \arg\max_{\theta \in \Theta} \mathcal{L}_S(\theta)$$

In the case of normal distributions, it is easily checked that we have $\widehat{\theta} = (\widehat{\mu}, \widehat{\sigma}^2)$, where $\widehat{\mu}$ and $\widehat{\sigma}^2$ are the *sample mean* and *sample variance* defined by:

$$\widehat{\mu} = \frac{1}{|S|} \sum_i s_i \qquad\qquad \widehat{\sigma}^2 = \frac{1}{|S|} \sum_i (s_i - \widehat{\mu})^2$$

Hence, we take $v = \widehat{\mu}$ as our current guess of the mean of variable X.

Example 4. Assume that we ask 3 users to evaluate sound samples compressed with ratio 3, and obtain the grades 3, 5, and 7. This means that our sample mean and variance for variable Q_3 are respectively $\widehat{\mu_3} = 5$ and $\widehat{\sigma_3}^2 = 8/3$.

Estimating the error. How to estimate the loss of our prediction $\widehat{\mu}$? Because the true value is unknown, we estimate the loss by assuming that our current guess $\widehat{\theta}$ is correct, and finding out what its expected error is. We do this by examining the

[4] Note that likelihood cannot, however, be seen as a probability distribution on Θ.

range of samples that could have been obtained instead of S under the assumed distribution and computing the loss of the MLE obtained from them.

By the central limit theorem, the distribution of the mean of N samples of $\mathcal{N}(\widehat{\mu}, \widehat{\sigma}^2)$ can be approximated by $\mathcal{N}(\widehat{\mu}, \widehat{\sigma}^2/N)$. Hence, under the assumption that $\widehat{\theta}$ is correct, we can define the average error obtained through the MLE method from $|S|$ samples, as follows:

$$E(\widehat{\theta}, |S|) = \int_{x \in \mathbb{R}} \Pr_{(\widehat{\mu}, \widehat{\sigma}^2/|S|)}(x) L_{\widehat{\mu}}(x) \, \mathrm{d}x$$

This integral can be numerically approximated by sampling.

Example 5. The estimated error for Q_3 under the samples S_3 of the previous example is the probability that the sample mean, distributed according to $\mathcal{N}(5, (8/3) \cdot (1/3))$, is above threshold $\tau = 6$ (as the loss is then 1, and is 0 otherwise, relative to our estimate $\widehat{\mu_3} = 5$). Numerically we have $E(\widehat{\theta}_3, |S_3|) = 0.144$.

Estimating the error decrease. Now that we can estimate the parameter of a distribution from the samples, and the expected error according to this parameter, we can easily devise an estimation of how this error may decrease when an additional sample is requested from variable X.

Let us assume that we obtain a new sample of X with value x, and call S' the $|S| + 1$ samples obtained by adding x to S. Call $\widehat{\theta}'$ the MLE obtained by maximizing $\mathcal{L}_{S'}$, and define the *error decrease* as $D(S, x) = E(\widehat{\theta}, |S|) - E(\widehat{\theta}', |S| + 1)$. This gives us an estimation of how error decreases for one more sample with value x. Of course, we cannot know if we would indeed obtain value x, but we can compute its probability according to our current hypothesis for the underlying distribution, namely $\mathcal{N}(\widehat{\mu}, \widehat{\sigma}^2)$. We therefore define the *expected error decrease*:

$$D(S) = \int_x \Pr_{\widehat{\theta}}(x) D(S, x) \mathrm{d}x$$

This is our estimate of the expected loss reduction when sampling this variable.

Example 6. If we obtain one additional sample of 5 for Q_3 (yielding S'_3), the estimated mean $\widehat{\mu}_3 = 5$ is unchanged, but the estimated variance decreases to 2 so the estimated error under the new MLE $\widehat{\theta}'_3$ becomes $E(\widehat{\theta}'_3, |S'_3|) = 0.079$. We estimate the expected error decrease by averaging the decrease under possible additional samples drawn from our estimated distribution $\mathcal{N}(\widehat{\theta}_3)$ for Q_3.

Multiple variables. With the above method, we can compute the expected error decrease of each variable, and sample the one whose expected error decrease is highest. It is easy to see that this greedy approach is optimal in terms of reducing the expected error over any fixed number of requests, as samples for one variable do not change the estimated parameter or expected error of other variables.

4 With order constraints

Under order constraints, the problem is more challenging. Though the loss function remains a sum of loss functions over individual parameters, it is not possible anymore to manage variables separately, because information obtained for one variable gives us additional information about the other variables. Reconsidering

our running example, under the objective of identifying the lossiest compression ratio with average quality at least τ, it makes little sense to consider the results of every variable independently, and we should examine the results globally to locate where the decreasing sequence of qualities intersects the threshold τ. The challenge is how to formalize such a global strategy, under general constraints.

To this end, we propose a greedy strategy inspired by that of the previous section, but integrating the order constraints and considering the variables globally rather than in isolation. Because additional samples on one variable give us information about other variables, such a greedy approach is no longer guaranteed to be optimal over multiple draws. Because of space constraints, we only sketch the principles of our initial approach; we plan to study this further and examine possible alternative approaches in future work.

We consider the parameter space $\Theta = (\mathbb{R} \times \mathbb{R}_+)^n$, covering all parameters of all random variables simultaneously, and we define the likelihood of $\boldsymbol{\theta} \in \Theta$ (with $\theta_i = (\mu_i, \sigma_i^2)$) as a function of $\mathcal{S} = (S_1, \ldots, S_n)$, the set of all samples for all variables, using the fact that all draws are still independent. We exclude parameters which violate order constraints by defining the likelihood as follows:

$$\mathcal{L}_{\mathcal{S}}(\boldsymbol{\theta}) = \begin{cases} \prod_i \prod_{s \in S_i} \Pr_{\theta_i}(s) & \text{if } E \cdot \boldsymbol{\theta} \leq (0) \\ 0 & \text{otherwise} \end{cases}$$

The main problem is now to determine the maximum likelihood estimator for $\boldsymbol{\theta}$ by maximizing this expression. We next propose a possible approach to the problem, and the challenges yet to be resolved.

Estimating the means. We propose to maximize the expression as a function of the means $\boldsymbol{\mu}$, while making the assumption that the variances are the sample variances $\widehat{\boldsymbol{\sigma}}$ for every individual variable. Under this approximation, the maximization problem can be rewritten as maximizing a quadratic expression with a positive definite matrix under the inequalities E. Such a problem is tractable [5], so we can solve it and obtain a set of candidate means $\mathbf{v}$ for the underlying distributions. Technical details are omitted for lack of space.

Example 7. Assume that we have obtained the same number of samples for Q_1, Q_2 and Q_3, that their sample variances are equal ($\widehat{\sigma_1} = \widehat{\sigma_2} = \widehat{\sigma_3}$), and that the sample means are $\widehat{\mu_1} = 9$, $\widehat{\mu_2} = 7$, and $\widehat{\mu_3} = 8$. Observe that we have $\widehat{\mu_2} < \widehat{\mu_3}$ even though we know that $q_3 \leq q_2$. In this specific setting, our estimation of the means is the solution $\mathbf{v}$ of a quadratic programming problem amounting to minimizing the sum of squares $\sum_i (v_i - \widehat{\mu_i})^2$ subject to the inequalities: its solution is $v_1 = 9$, $v_2 = 7.5$, and $v_3 = 7.5$.

Estimating the variances. We have computed the MLE estimator for the means of the distributions subject to the inequality constraints, up to the approximation of substituting the individual sample variances instead of integrating them in the maximization problem. Since the estimations of the means and variances are inter-dependent, we may now need to reestimate the variances.

Example 8. Assume that we have samples $S_2 = \{0.1, 0.2\}$ for Q_2, and numerous samples for Q_1 and Q_3 which convince us that $v_1 = 9$ and $v_3 = 8.5$ are very

good estimates for q_1 and q_3. We know that we must have $8.5 \leq v_2 \leq 9$ (we will probably choose $v_2 = 8.5$ given S_2), but then our estimation of the variance of Q_2 should be much higher than the sample variance $\widehat{\sigma}_2^2$ of S_2 in isolation.

We estimate the variance of each X_i under the computed means $\mathbf{v}$ (and thus estimate the complete parameter $\boldsymbol{\theta}$) as the sample variance relative to the computed mean v_i of X_i (instead of relative to the sample mean). The solution thus obtained may not be optimal, as we have fixed and optimized the means and variances separately rather than simultaneously. Estimating how much this approach deviates from the true solution is a challenge for future work.

Estimating the error and error decrease. The overall method now follows Section 3 except that we follow the above[5] to fit a family of distributions to the variables.

5 Interpolation

In some real-life scenarios, we may have a very large number of questions to ask the crowd; for instance, the number of possible compression ratios may be very high, almost continuous. In such cases, we may have many variables X_i with no samples at all: those variables thus do not appear in the optimization problem, so that we know nothing about them (except that they satisfy the order constraints). However, we could then perform *interpolation* to estimate more precisely a large proportion of the variables with a limited number of questions to the crowd.

In the general case where E is an arbitrary set of inequalities, it is hard to define how to interpolate a value for a variable with no samples. We leave this general question to future work, and only focus on the case where E expresses the total order $\mu_1 \geq \cdots \geq \mu_n$. For simplicity, up to renumbering indices, we assume that we have a model for X_1 and X_n, namely (μ_1, σ_1^2) and (μ_n, σ_n^2), and that we wish to derive a model for X_k, $1 \leq k \leq n$, for which we have no samples.

Example 9. If we estimate $v_1 = 8$ and $v_5 = 4$, our best guess for q_3 in the absence of samples should be $v_3 = 6$. Likewise, our best guess for q_4 should be $v_4 = 5$.

Interpolating the mean. We interpolate the mean μ_k by a linear interpolation between μ_1 and μ_n according to the rank k, as presented in Example 9

Interpolating the variance. We want to interpolate σ_k^2 by combining both the variances of X_1 and X_n, and the uncertainty arising from the interpolation itself: the further away k is from 1 and n, the least certain we are about μ_k.

To do so, we consider that μ_k has been chosen by picking $n-2$ random uniform values between μ_1 and μ_n (the means $\mu_2, \ldots, \mu_{n-1}$), sorting them, and choosing the $(k-1)$-th value to be μ_k. Now, this means that μ_k is the $(k-1)$-th *order statistic* of $n-2$ uniform and independent random variables in $[\mu_1, \mu_n]$, so that it follows a *beta distribution* [3] whose variance has a closed form.

Example 10. Pursuing Example 9, for $\mu_5 = 8$ and $\mu_9 = 4$, we estimate the variance on μ_7 to be $4/5$ for this outcome (that of the adequate beta distribution).

[5] Note that this also changes the way of fitting distributions when computing the error decrease under possible additional samples.

We can thus estimate a variance for X_k for fixed values μ_1 and μ_n: those values are unknown, but can be sampled according to our model for X_1 and X_n to yield an overall variance for X_k. We omit details for lack of space, and leave to future work the study of other possible interpolation methods for variance.

6 Conclusion and perspectives

In this paper, we have studied the problem of learning numerical values from the crowd, leveraging ordering constraints on those values to mitigate the uncertainty on crowd answers. We have presented an abstract framework inspired by [2] ignoring the order constraints, and presented an approximate method to take those constraints into account, along with a way to interpolate values for yet unsampled variables. We have identified further challenges to be explored.

Our main direction for future work is to study more carefully the approximations and design choices that we made, and to implement our approach to evaluate its effectiveness. We plan to evaluate, over various datasets and objectives, the importance of accounting for order constraints and performing interpolation, and compare our approach to round-robin or random baselines, as well as ad-hoc strategies for specific scenarios such as total orders.

Acknowledgements. This work has been partially funded by the European Research Council under the FP7, ERC grant MoDaS, agreement 291071, and by the Israel Ministry of Science.

References

1. A. Amarilli, Y. Amsterdamer, and T. Milo. On the complexity of mining itemsets from the crowd using taxonomies. In *Proc. ICDT (to appear)*, Athens, 2014.
2. Y. Amsterdamer, Y. Grossman, T. Milo, and P. Senellart. Crowd mining. In *Proc. SIGMOD*, pages 241–252, New York, USA, June 2013.
3. H. A. David and H. N. Nagaraja. *Order Statistics*, chapter 2, page 14. Wiley, 2003.
4. R. M. Karp and R. Kleinberg. Noisy binary search and its applications. In *Proc. 18th ACM-SIAM Symp. on Discrete Alg.*, 2007.
5. M. K. Kozlov, S. P. Tarasov, and L. G. Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Comp. Math. and Math. Phys.*, 20(5), 1980.
6. Amazon Mechanical Turk. `https://www.mturk.com/`.
7. A. Parameswaran, H. Garcia-Molina, H. Park, N. Polyzotis, A. Ramesh, and J. Widom. Crowdscreen: Algorithms for filtering data with humans. In *Proc. SIGMOD*, 2012.
8. A. Parameswaran, A. Sarma, H. Garcia-Molina, N. Polyzotis, and J. Widom. Human-assisted graph search: it's okay to ask questions. *Proc. VLDB*, 4(5), 2011.
9. D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proc. AAAI/IAAI*, 2000.
10. E. Triantaphyllou. *Data Mining and Knowledge Discovery by Logic-Based Methods*, chapter 10. Springer, 2010.
11. B. Trushkowsky, T. Kraska, M. J. Franklin, and P. Sarkar. Crowdsourced enumeration queries. In *Proc. ICDE*, 2013.
12. X. Yang, R. Cheng, L. Mo, B. Kao, and D. Cheung. On incentive-based tagging. In *Proc. ICDE*, 2013.